

VIRTIO HALS

ALEX BENNÉE

GLOBAL SOFTWARE TECHNOLOGY
SUMMIT 2022

INTRODUCTION

WHO AM I

- Alex Bennée
- Senior Virtualisation Engineer @ Linaro
- [Project Stratos](#) Tech Lead
- IRC: stsquad/ajb-linaro

PROJECT STRATOS

- High Performance VirtIO
- Standardised hypercalls
- Boot orchestration
- VM's with safety islands

TOPICS

- Intro to VirtIO
- Deployment scenarios
- Project Stratos Implementations
- Limitations and Further work

VIRTIO INTRODUCTION

HISTORY

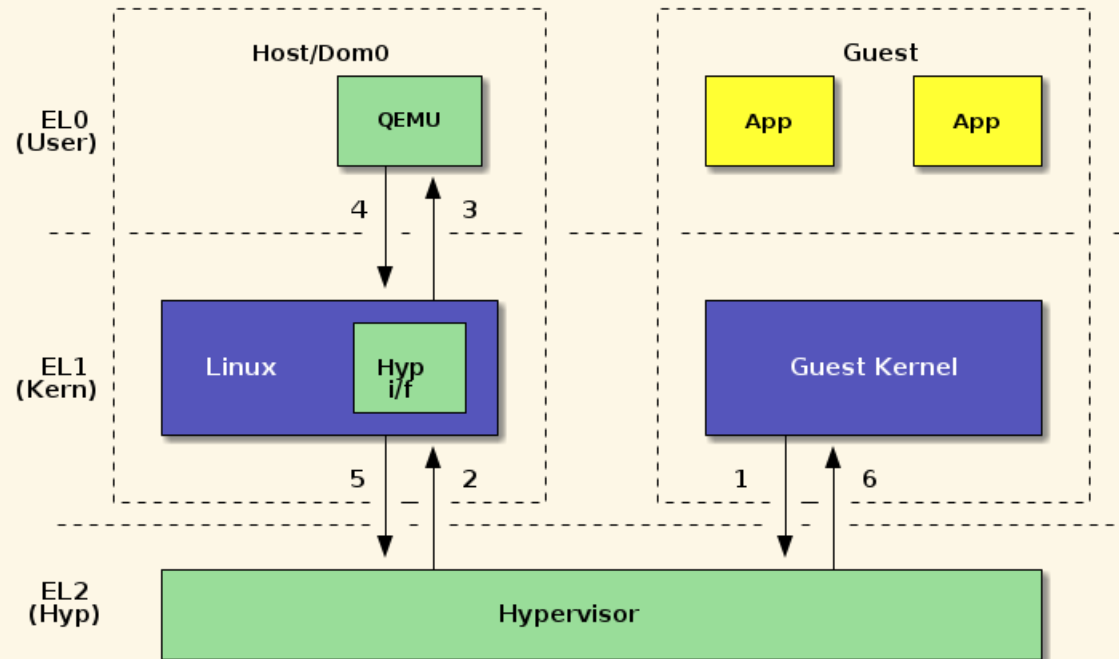
- Merged in [v2.6.24](#) (2007)
- OASIS v1.0 (2016)
- OASIS v1.1 (2019)
- [OASIS v1.2](#) (2022)

TERMINOLOGY

- Guest OS
 - VirtIO Front-end
 - Device Driver
- Transport
 - PCI, MMIO, Channel I/O (s390x)
- Host OS
 - VirtIO Back-end

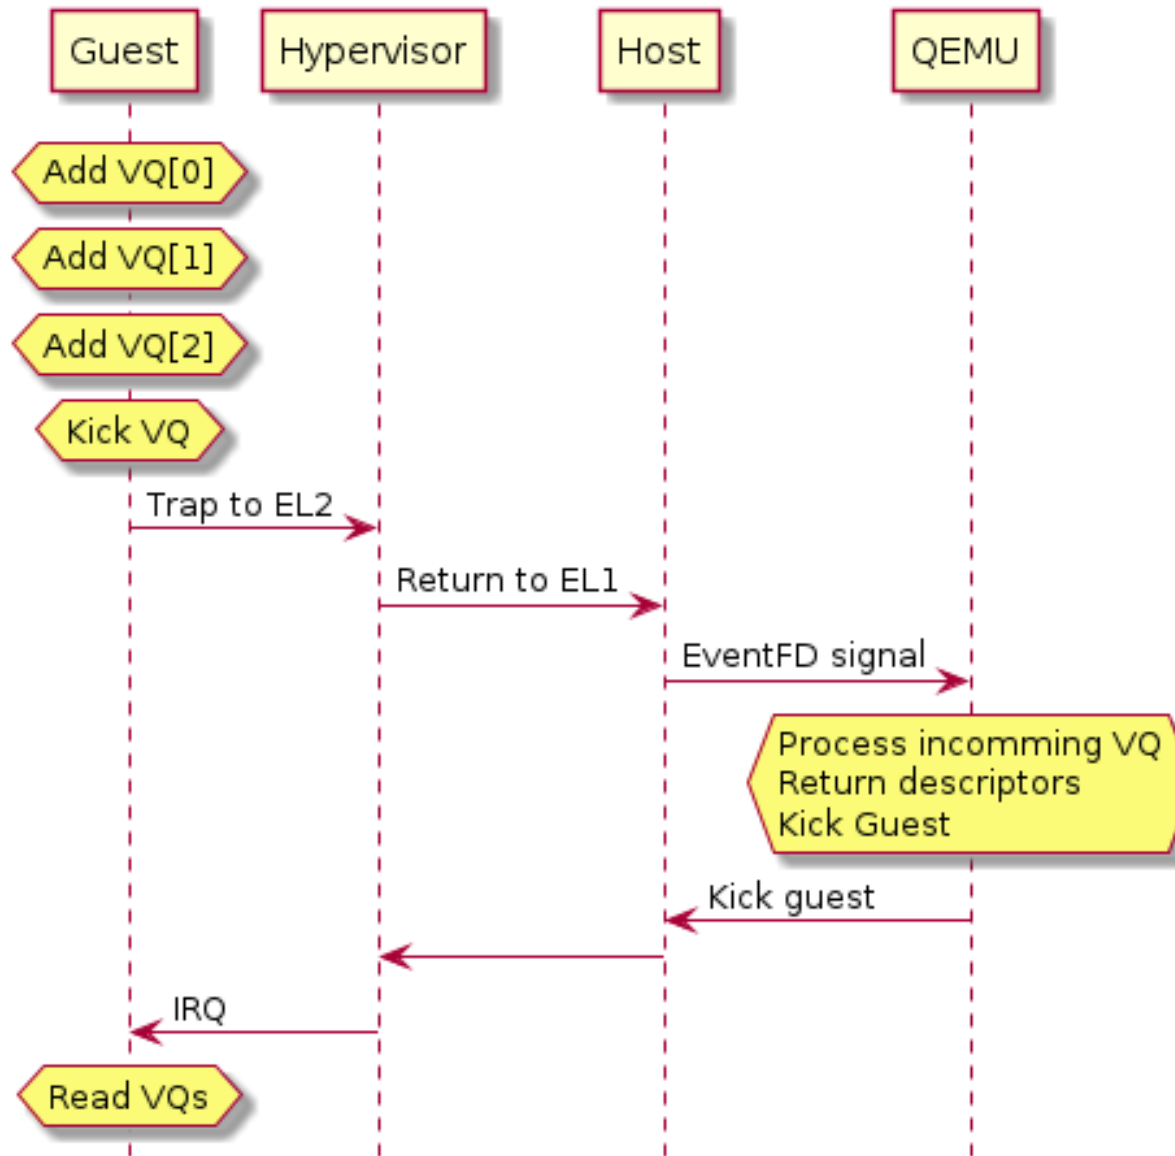
HOW IT WORKS

TYPICAL DEVICE EMULATION



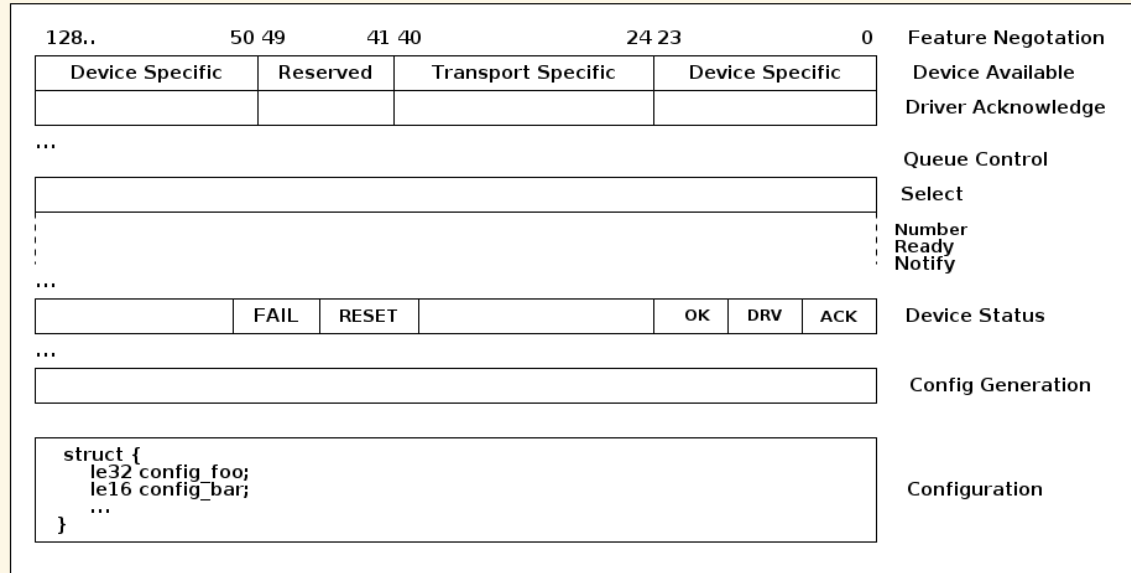
TYPICAL VIRTIO TRANSACTION

Virtio Sequence Diagram




VIRTIO DEVICES

Device Registers



Virt Queues



DEVICE TYPES

virtio-net virtio-block virtio-console

DEVICE TYPES

virtio-net	virtio-block	virtio-console
------------	--------------	----------------

virtio-scsi	virtio-rng	virtio-balloon
-------------	------------	----------------

DEVICE TYPES

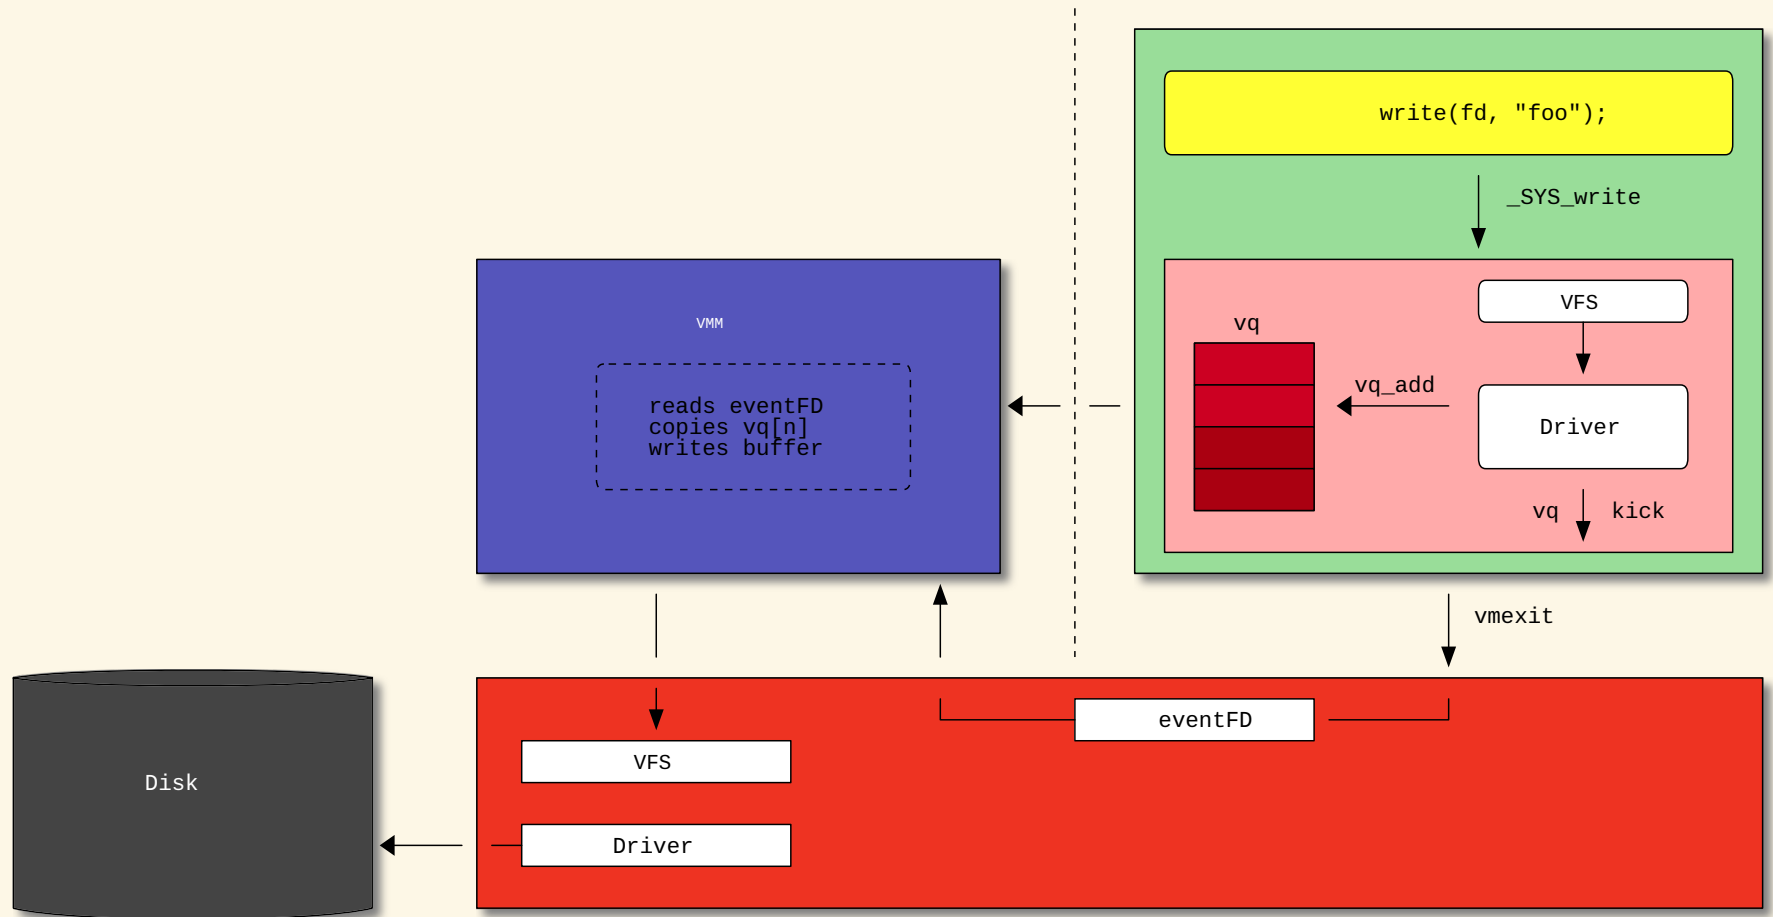
virtio-net	virtio-block	virtio-console
------------	--------------	----------------

virtio-scsi	virtio-rng	virtio-balloon
-------------	------------	----------------

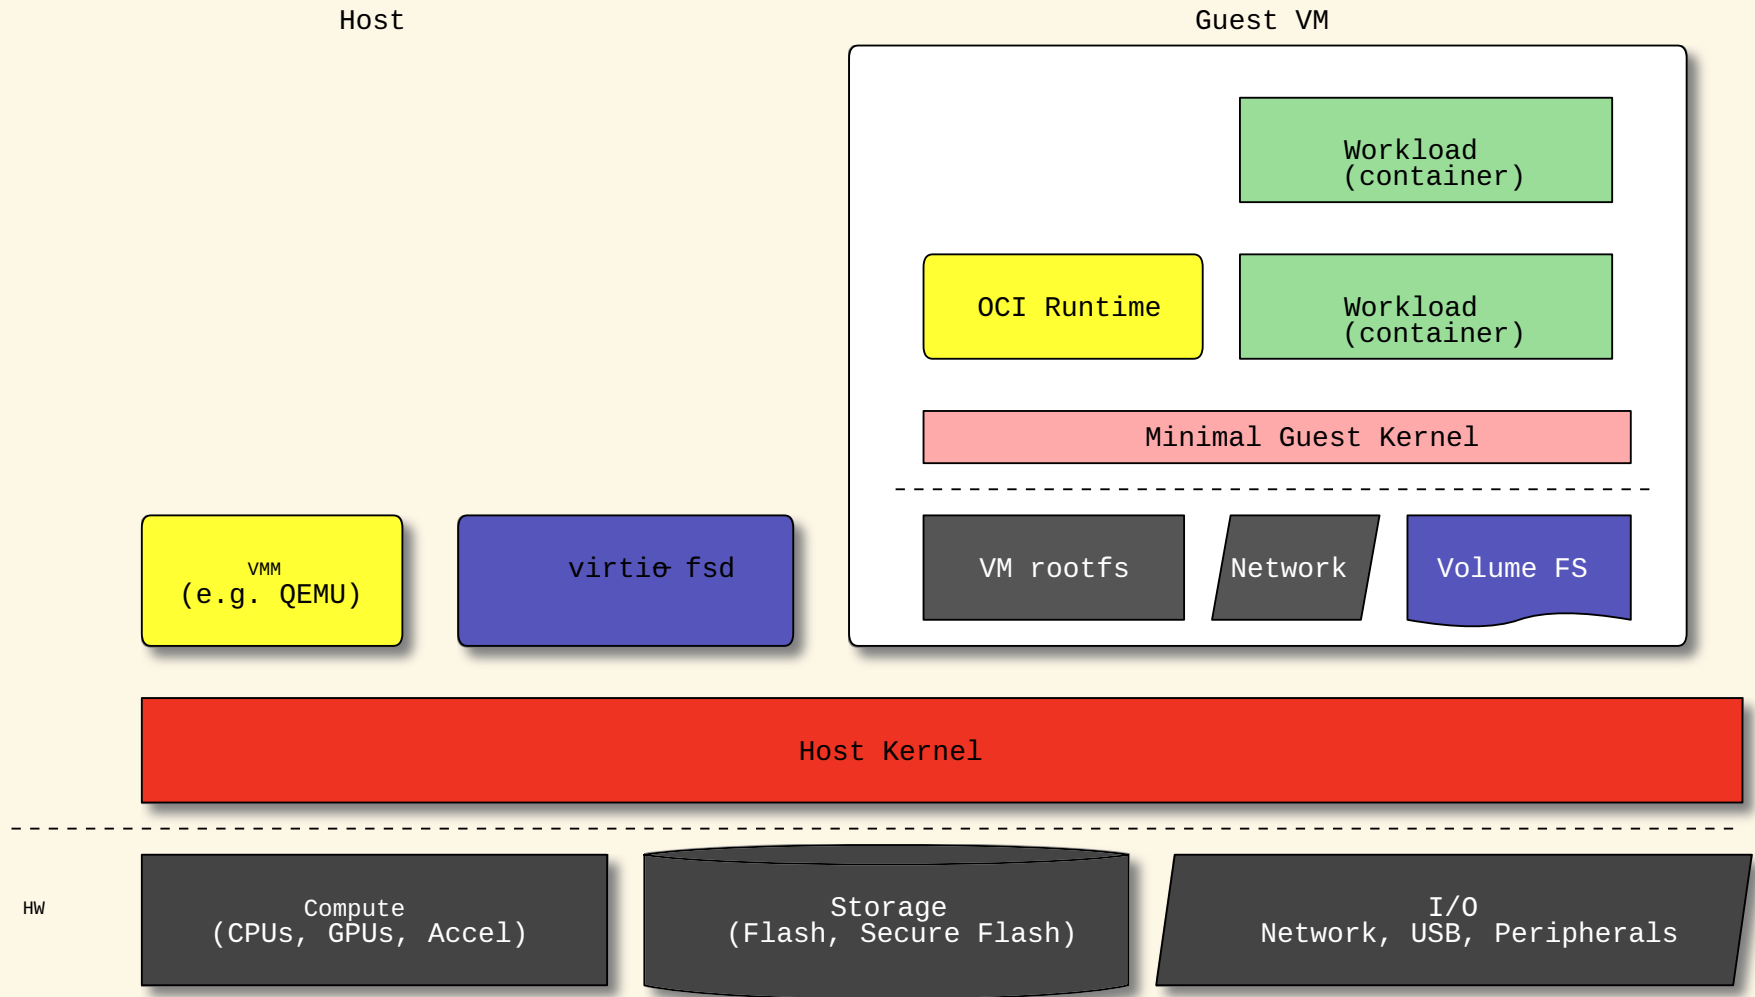
virtiofs	vsock
----------	-------

DEPLOYMENTS

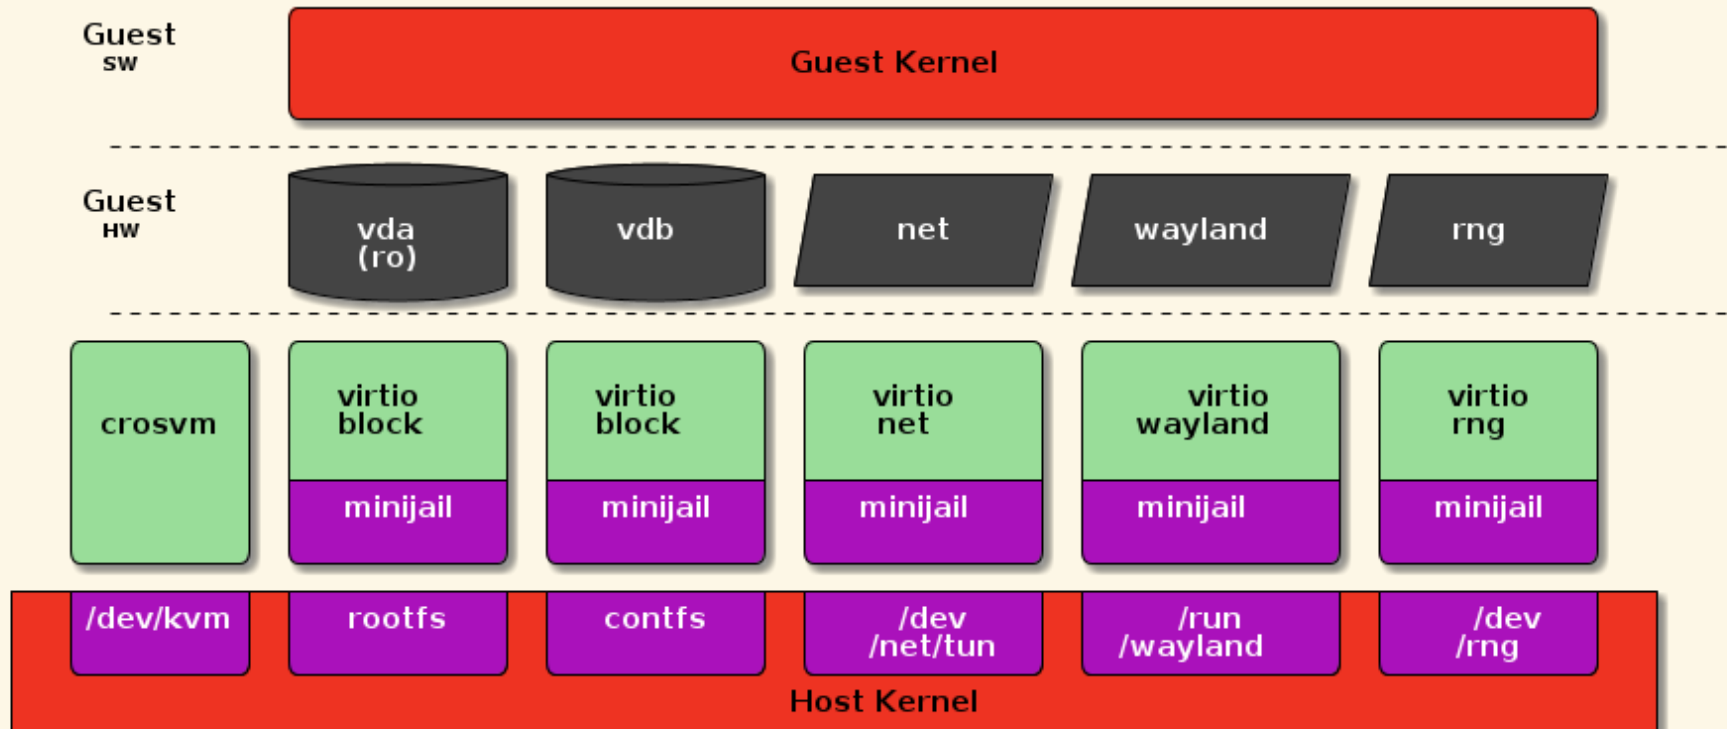
TRADITIONAL



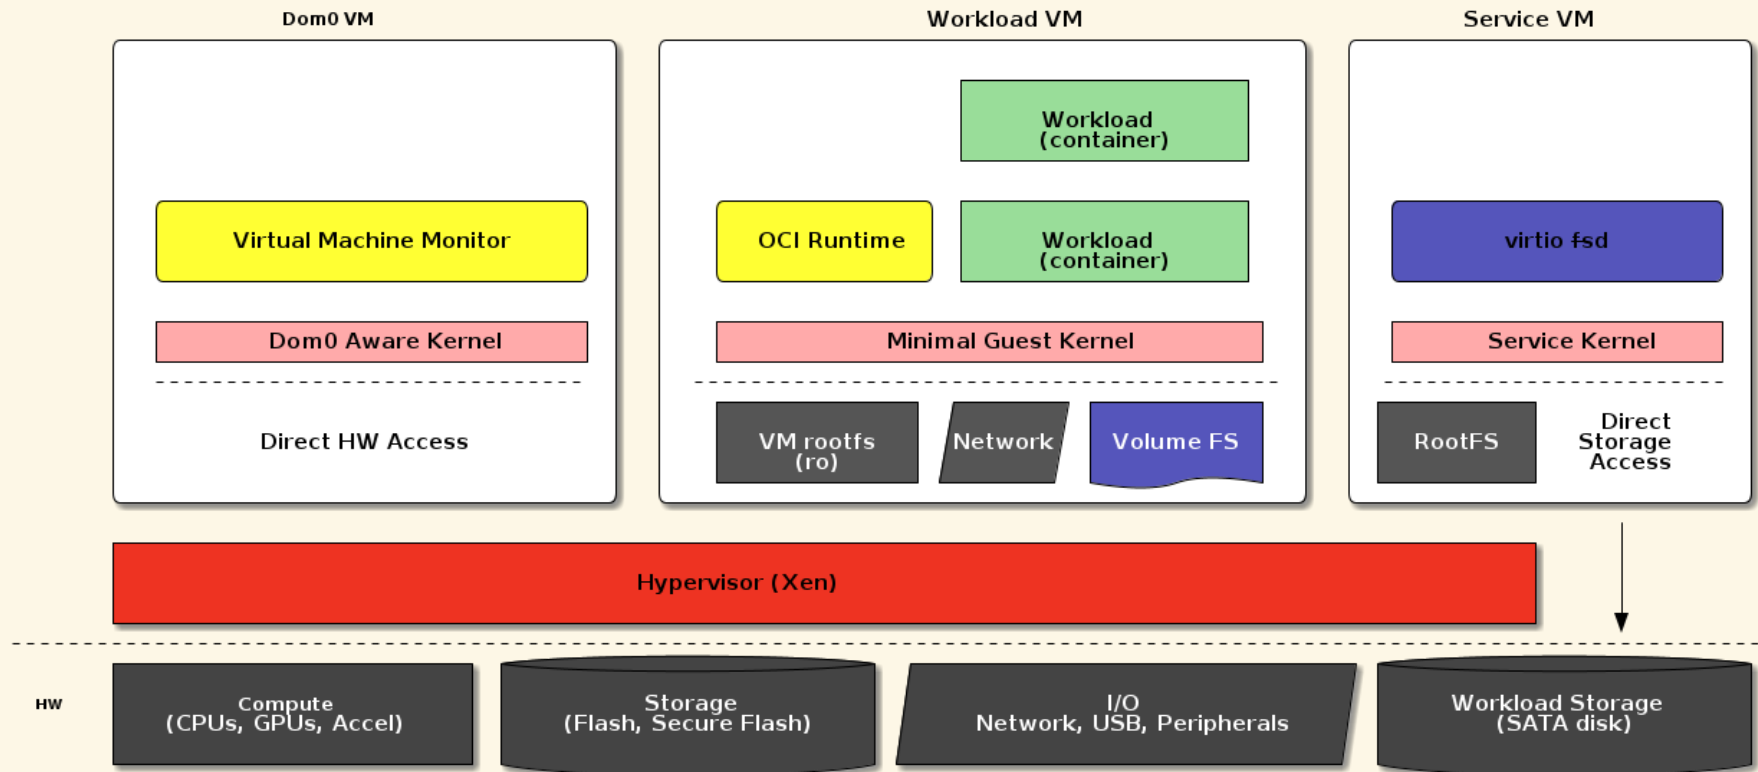
SEPARATE BACKENDS



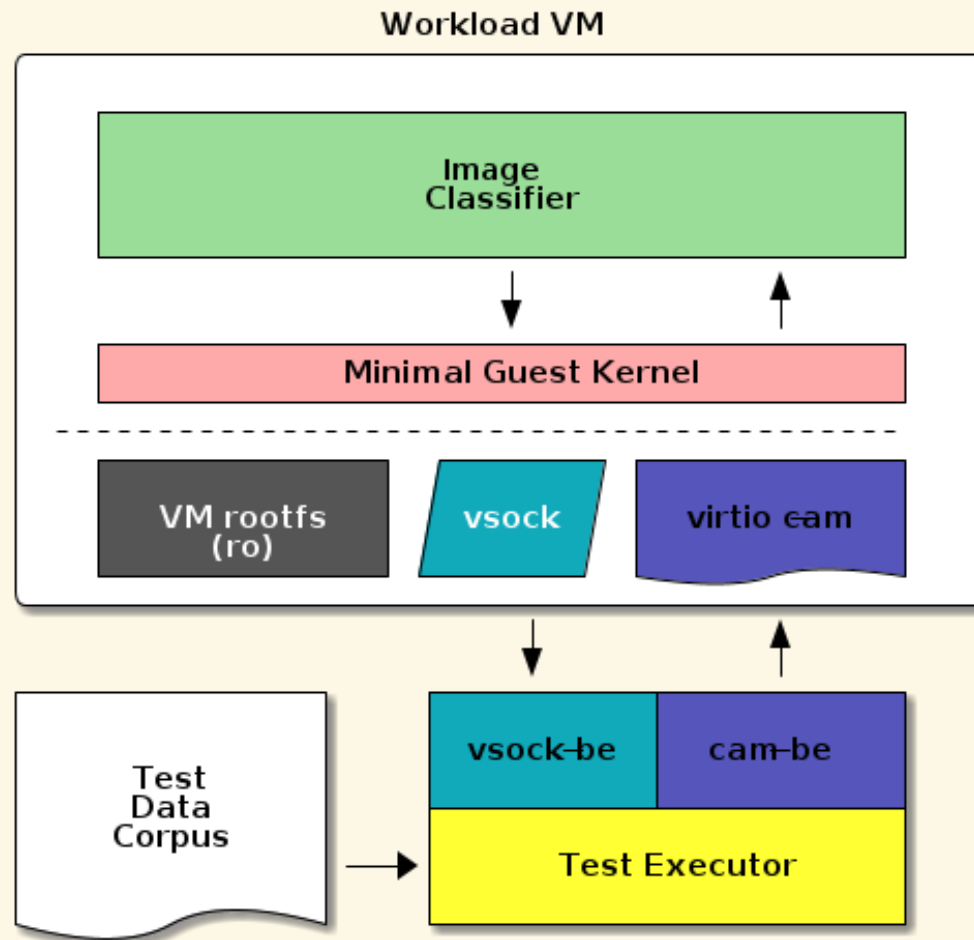
CROSVN



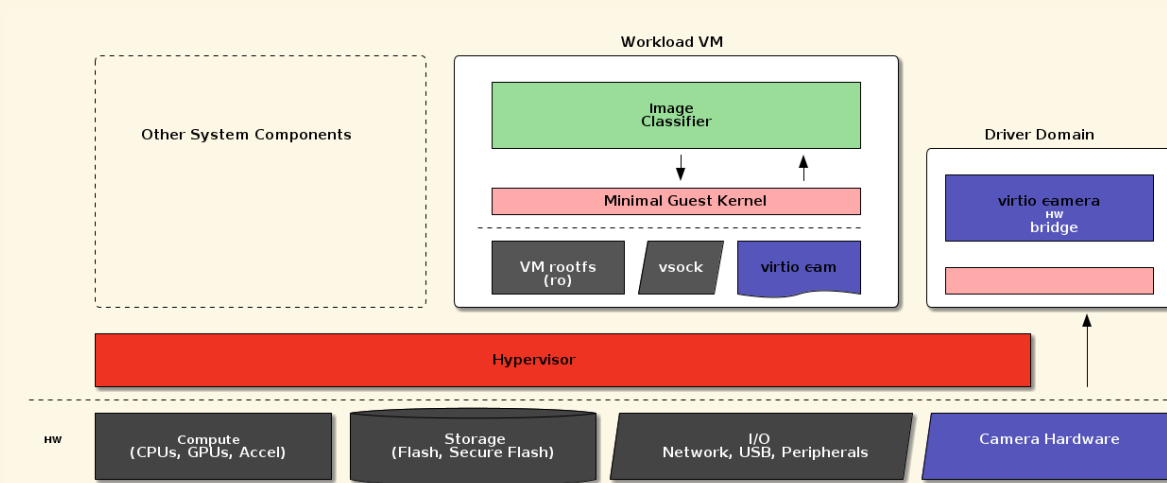
DRIVER DOMAINS



CLOUD NATIVE



HW DEPLOYMENT



PROJECT STRATOS IMPLEMENTATIONS

NEW VIRTIO-DEVICES

- virtio-scmi
- virtio-rpmb
- virtio-rng
- virtio-gpio
- virtio-i2c

SCMI

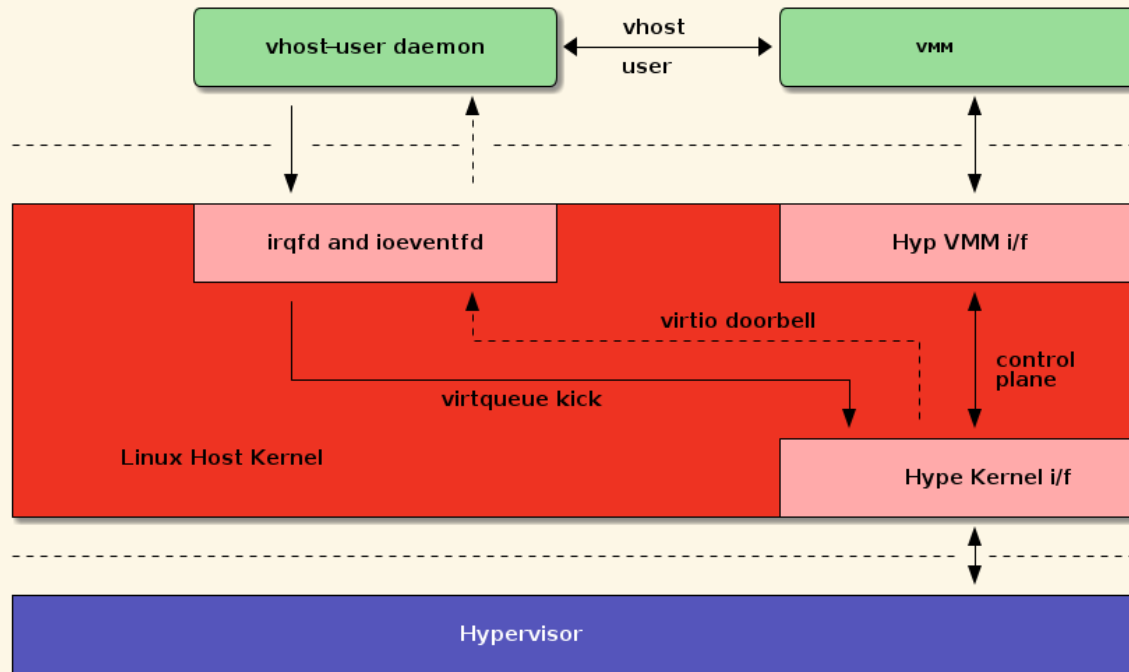
- System Control and Management Interface
- Power, performance and system management
- Often resides in firmware
- Access through Secure Monitor calls
 - SMC or HVC

VIRTIO-SCMI

VHOST-USER DAEMONS

- userspace daemons
- using rust-vmm
- portable* between hypervisors

VHOST-USER



RUST-VMM

- Component approach
 - Run loops (vm control)
 - Memory management
 - VirtIO devices
 - Other device emulation
- crates.io

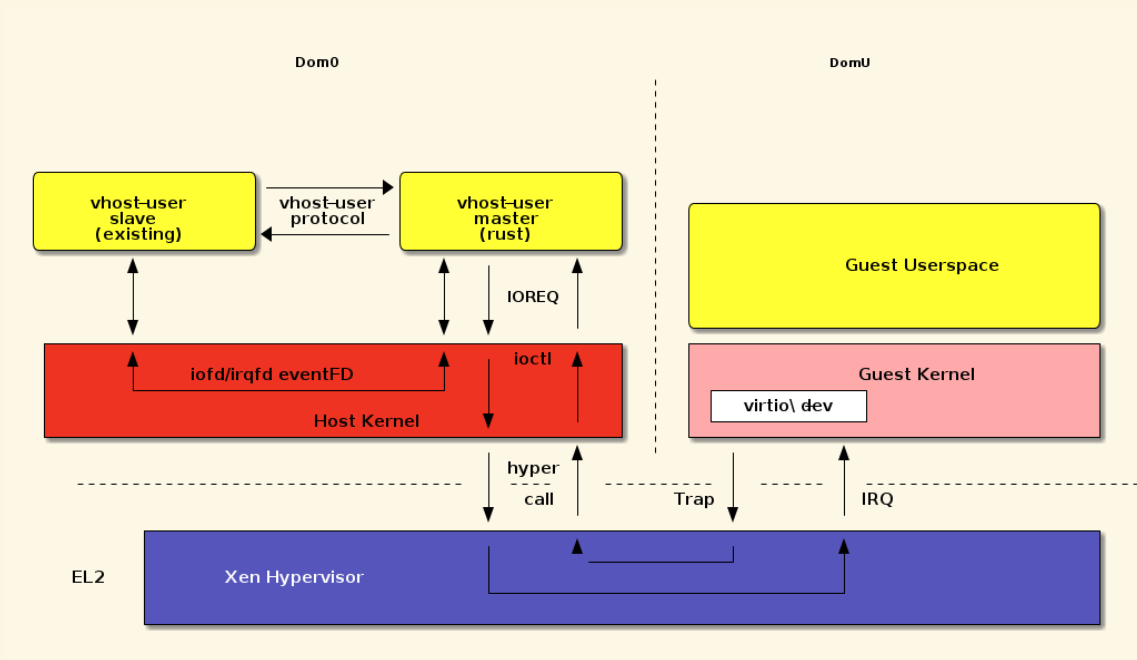
VHOST-DEVICE CRATE

- Device emulation
 - i2c
 - rng
 - gpio
- RFCs
 - vsock
 - virtio-scsi

OXERUN CRATE

- Xen Hypercalls
- ioctl via kernel
- direct bare metal

XEN-VHOST-MASTER



FUTURE WORK

IMPROVE XEN API

- More standard mmap
- direct irqfd/eventfd routing

ISOLATION

- just started exploring
- fat virtqueue
- iommu/grants vs regions

BARE METAL RUST

- re-use exiting VirtIO logic
- but without POSIX layer

JOIN US!

- [Project Stratos Homepage](#)
- [Mailing List](#)
- Fortnightly Sync

QUESTIONS?

EXTRA SLIDES

VHOST DIAGRAM

